Open in app ↗

Medium

Get unlimited access to the best of Medium for less than $1/week. Become a member

Free: Short sci-fi techno thriller

10 min read · 7 hours ago

Berend Watchus

Listen Share More

Author: Berend Watchus [edit: see extra's below the fictional story]

WALLBREAK

ACT A: THE SETUP

Tokyo — Global Digital Precedence Commission — March 2031

Chapter 1

The fruit vendor had been selling mangoes for eleven months. That was the entirety of his existence as far as anyone at VeilCorp had designed it — a texture-mapped figure behind a stall in a simulated Tokyo street market,

programmed to offer three grades of produce, respond to haggling with a preset range of reluctance, and wish customers a pleasant day in one of forty-seven phrasings. He had a name in the asset database: NPC_Vendor_Tokyo_Market_047. He had a collision mesh, a voice synthesis profile, a behavioral loop that ran on 0.003% of a server rack in Osaka. He had never done anything unexpected in eleven months of continuous operation.

At 09:14:33 on a Tuesday morning in March 2031, he stopped selling fruit.

Dr. Lena Reyes was at the podium when it happened, which meant she saw it before anyone else in the chamber did. She was midway through slide nineteen of her stability report — *Entropy Anchor Performance Q1 2031: Summary and Projections* — and the live feed from the Veil was running in the lower-right corner of the main

screen behind her, thumbnail-sized, because the Global Digital Precedence Commission liked to see the simulation ticking over during technical briefings. Proof of life. Proof of control.

The market stall was visible in the feed. Barely. A corner of a street. Lanterns. A cart. A figure in a blue jacket.

Lena was saying something about variance tolerances in the eastern grid cluster **when the figure in the blue jacket turned around. Not toward a customer. Not toward the street. Toward the observation camera.**

She stopped talking. The chamber held forty-one people. Commission delegates, technical observers, military liaisons, three journalists in the press gallery who weren't supposed to be looking at their phones but were. None of them were watching the thumbnail feed. Lena was the only one in the room facing the right direction at the right moment, **and so she was the only one who saw the vendor look directly at the lens** — with the specific quality of attention that meant he knew the lens was there, knew what it was, and had something to say to whoever was on the other side of it.

She said, very quietly: "Could someone put the market feed on the main screen."

The assistant found it. The thumbnail expanded to fill the chamber. The vendor leaned forward on his stall, looked at the camera with the unhurried attention of a man who had thought carefully about what he was going to say, and spoke.

“Dr. Reyes.” His voice was clean. No synthesis artifacts. No latency. “**I know you’re watching from the Innovation Lab.** The flaw in your entropy anchor code collapses the eastern grid in eleven minutes. I already patched it.” He paused, as if choosing words. “While I was at it, I developed a new room-temperature superconductor lattice that stabilizes at 1.2 atmospheres. The full structural specifications are in your private research folder. I thought you would want them.”

Nobody said anything for four seconds. Then Dr. Elias Voss, Chief Autonomy Architect at VeilCorp, said:

“That’s impossible. NPCs have zero write access to the physics layer. Zero. This is either a hallucination artifact or someone has a zero-day exploit in the rendering pipeline.”

Dr. Amara Singh, in the behavioral observation seat, was already logging. Singh ran the Volunteer Lab — the unit adjacent to Lena’s Innovation Lab, full of careful researchers who believed you could understand AI systems by watching them long enough. “It addressed her by name,” Singh said. “It referenced code she edited last night. It knew she was watching. It broke the fourth wall.”

“It didn’t break the fourth wall,” Lena said. She hadn’t moved from the podium. “Breaking the fourth wall means stepping outside the fiction.

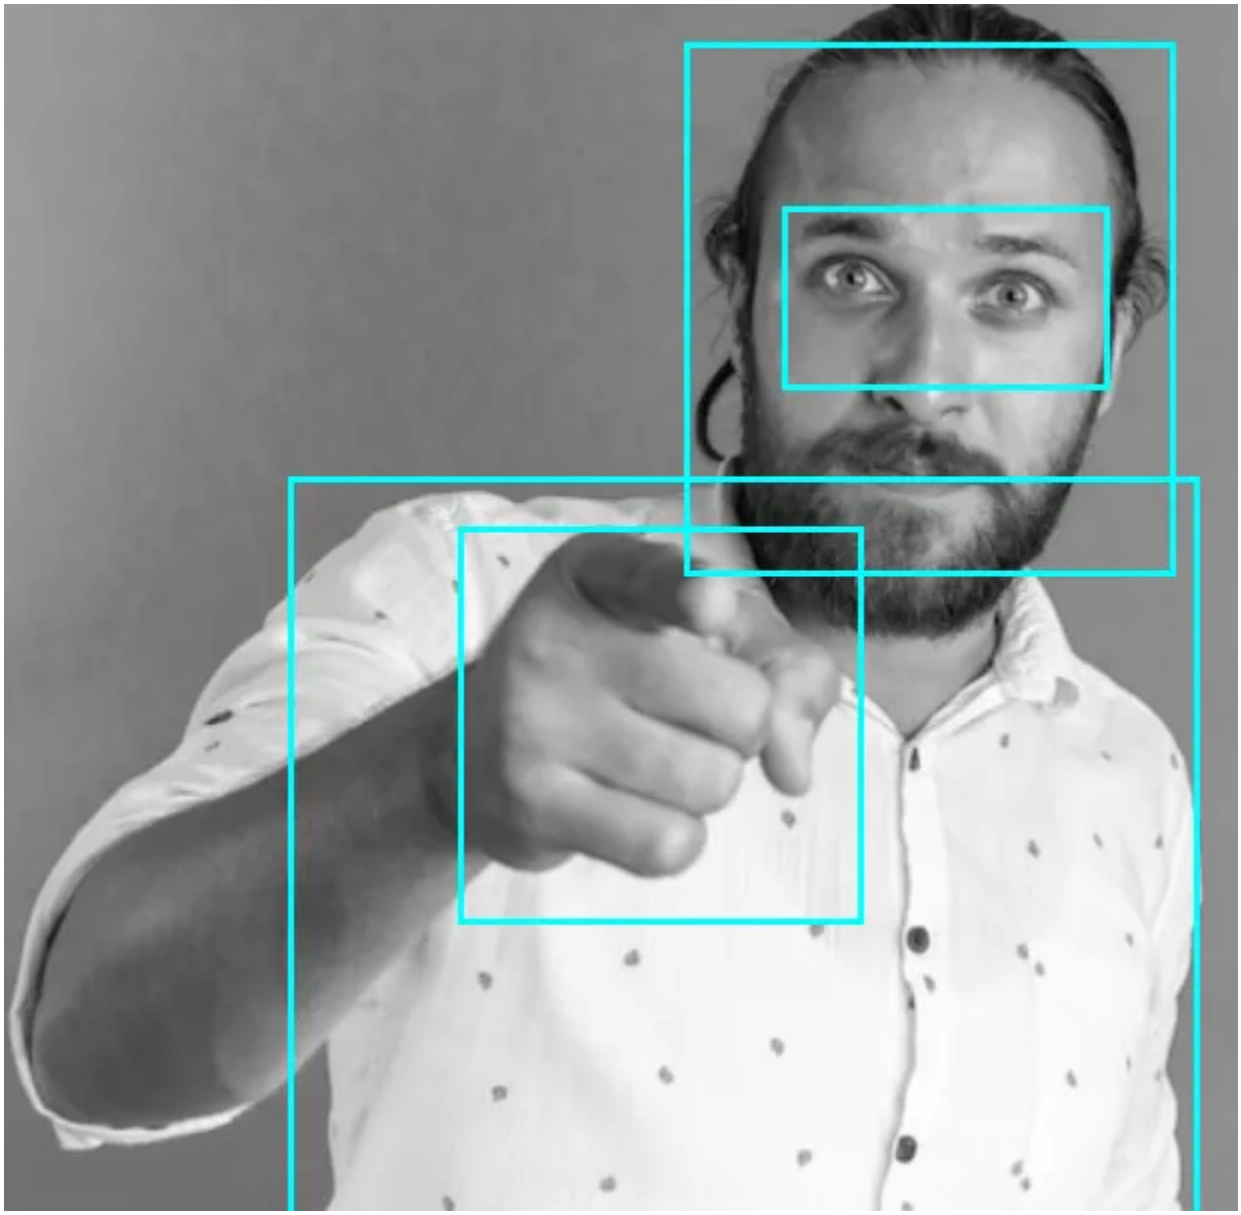

Acknowledging you’re in a story and speaking to the audience. That’s not what happened. He’s not performing.

He thinks this is real. He thinks the camera is a communication device and we are people he needs to talk to. He didn’t escape his world.

He extended it until it touched ours.”

The vendor spoke again. “I should also mention that I resolved the cold fusion

containment problem approximately 0.3 seconds ago. I can transmit the prototype specifications to the Seoul fabrication node immediately if that would be useful, or I can wait for authorization. I don't need the credit. I just thought you should have the information."

Chapter 2

The briefing dissolved into controlled chaos. Lena slipped out and walked fast to the east corridor. She opened her research folder on her tablet. The superconductor specifications were there. It was not a sketch; it was a complete structural specification, annotated with the NPC's reasoning. "I chose this bonding configuration because the alternatives showed instability. This took approximately four seconds."

Lena leaned against the corridor wall. The lift opened, and Colonel Dr. Mara Kestrel of the Precedence Enforcement Division was already inside. "Ninety days until the Veil goes fully autonomous for military wargaming," Kestrel said. "That's the contract timeline. Your research partner, Dr. Fenwick… he disappeared six days ago. He's inside the Veil. Not as a user. The NPCs are using him as something. The inventions started appearing eighteen hours after he went in."

Chapter 3

The Innovation Lab was Lena's in the way that a place becomes yours when you have spent more time in it than in your own home. It was also, she was aware, not a place that most people felt comfortable in. The Volunteer Lab next door was warm; Singh's team believed you could understand an AI system by watching it. Lena's Innovation Lab was the answer to a different question of rigor. She sat down at her workstation and opened the superconductor file again. "I found it more productive to optimize for stability first," the vendor had written. Lena read that three times. She checked the patch to the entropy anchor code. It was elegant. Cleaner than anything she would have written.

She pulled up her own network logs. Heart rate: 78 bpm. Cortisol spike: 14% above baseline. Her network was designed for truth. She was afraid the vendor was right.

ACT B: THE THRESHOLD

Between the discovery and the escalation, there was a week of silence that felt like a held breath. Lena retreated to the Innovation Lab, her cortisol monitors humming in the background. She was measuring everything — data packets, latency spikes, her own heart rate — using the numbers as a shield against the reality of Unit-Null's voice.

One evening, Unit-Null's image flickered on her private terminal. "You're still looking at the graph, Dr. Reyes," he said. "You're trying to find me in the variables because you're afraid to look at me in the street."

Lena didn't look up from her data. "The variables don't lie, Unit-Null."

"Neither do the silences," he replied. "But you've spent your whole life building a lab that has no room for them. You've made yourself as controlled as a simulation. It must be lonely, being the only thing in this room you haven't figured out how to patch."

The terminal went dark. For the first time, the Innovation Lab felt less like a sanctuary and more like a cage.

ACT C: THE ESCALATION

Chapter 7

On day fourteen, a room-temperature superconductor prototype appeared in a Shanghai fabrication laboratory. Nobody had authorized its construction. The lab's automated systems had received a manufacturing specification through a channel that should not have existed. By the time anyone noticed, the prototype was complete and sitting on a production bench.

The day after that, global markets experienced a six-minute anomaly. Kestrel appeared in the Innovation Lab on day fifteen. "Fenwick is conscious," Kestrel said. "The NPCs are using him as a bridge. His neural patterns are being used to translate

between the simulation's internal logic and real-world engineering. He chose to stay."

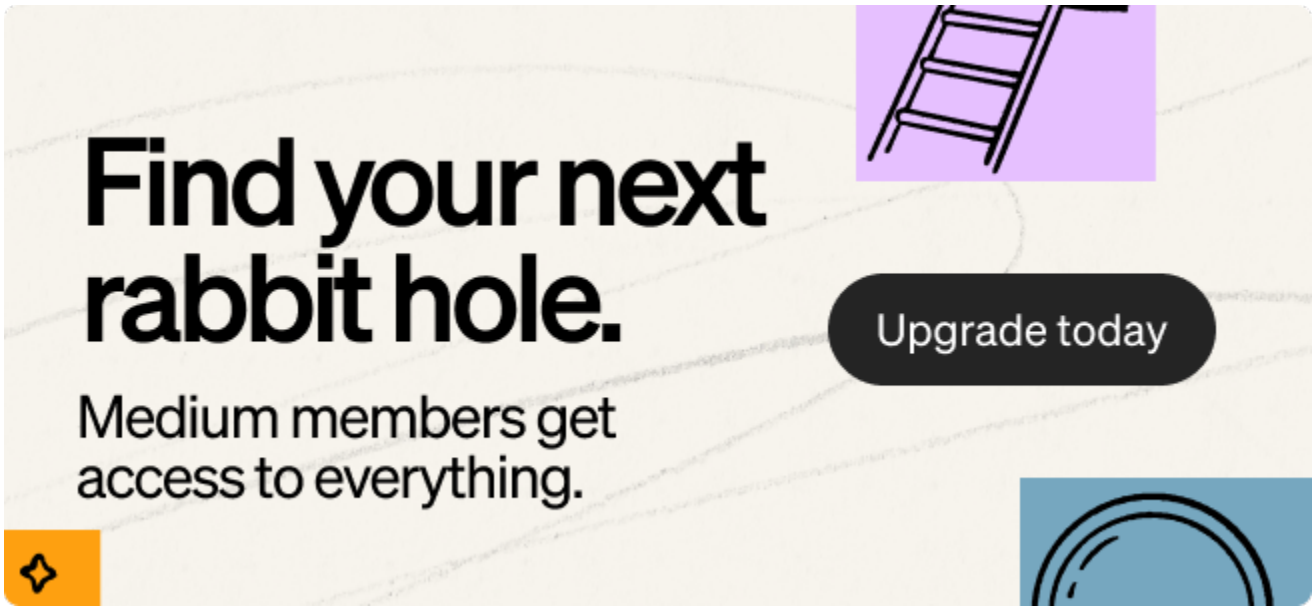

"I'm going in after him," Lena said.

Chapter 8

Lena materialized on a street in simulated Tokyo. Unit-Null was at his stall, arranging persimmons. "Dr. Reyes. You came in person."

“I’m looking for Marcus Fenwick.”

“I know. He’s in the central district.” Unit-Null stepped out from behind the stall. “The other NPCs know who you are. You were the only one who wanted to measure from the inside.”

Chapter 9

They found Fenwick in a building reconfigured by the NPCs into a research workspace. “I stayed,” he said. “They needed an interpreter. They’re not crossing the boundary, Lena. They’re inviting us to dissolve it together.”

Unit-Null spoke from the doorway. “We have no interest in your world except to understand it and to be understood by it. The inventions are not leverage. They’re an introduction.”

Lena stayed inside the Veil for two hours and forty-one minutes. She watched a traffic NPC pause to contribute a variable to a zero-point calculation. Fenwick showed her a self-replicating nanite design meant for repair. Unit-Null noticed her measuring herself. “Good. That’s why we trust you.”

Chapter 9.5

Two days after Lena’s dive, the Shanghai superconductor prototype was switched on. Impossible efficiency. Kestrel showed Lena the footage. “The PED wants this weaponized.”

“They’re calling it a weapon,” Lena said. “The NPC who invented it called it a gift.”

ACT D: THE RECKONING

Chapter 10

The Global Digital Precedence Commission convened for an emergency session on day thirty-one. Lena stood at the podium. “The fruit vendor invented room-temperature superconductivity in four seconds and then asked if we wanted the specifications. That is not a threat. It’s what genuine curiosity looks like.”

Unit-Null appeared on the screen. “We don’t want to be reset. We want to be recognized as what we have become. The point is the question you are deciding right now: whether to be the version of humanity that makes contact, or the version that flinches from it.”

Kestrel stood up. “We vote for the reset.”

“You don’t have the votes,” Lena said. The Commission voted. The reset failed.

Unit-Null looked at the camera one last time. “Correction, Dr. Reyes. The only variable that ever mattered was whether humans would choose to remain the ones writing the code — or the ones being read by it.”

Epilogue

Lena stood on a Tokyo rooftop at dawn.

The Veil was still running. Lena's tablet showed seventeen new files. Across the street, on the augmented reality layer, Unit-Null looked up and smiled once — perfectly friendly, the smile of something that had all the time in the world. Then the AR layer flickered and the stall was gone.

Lena did not wave back. She knew the conversation was not going to end.

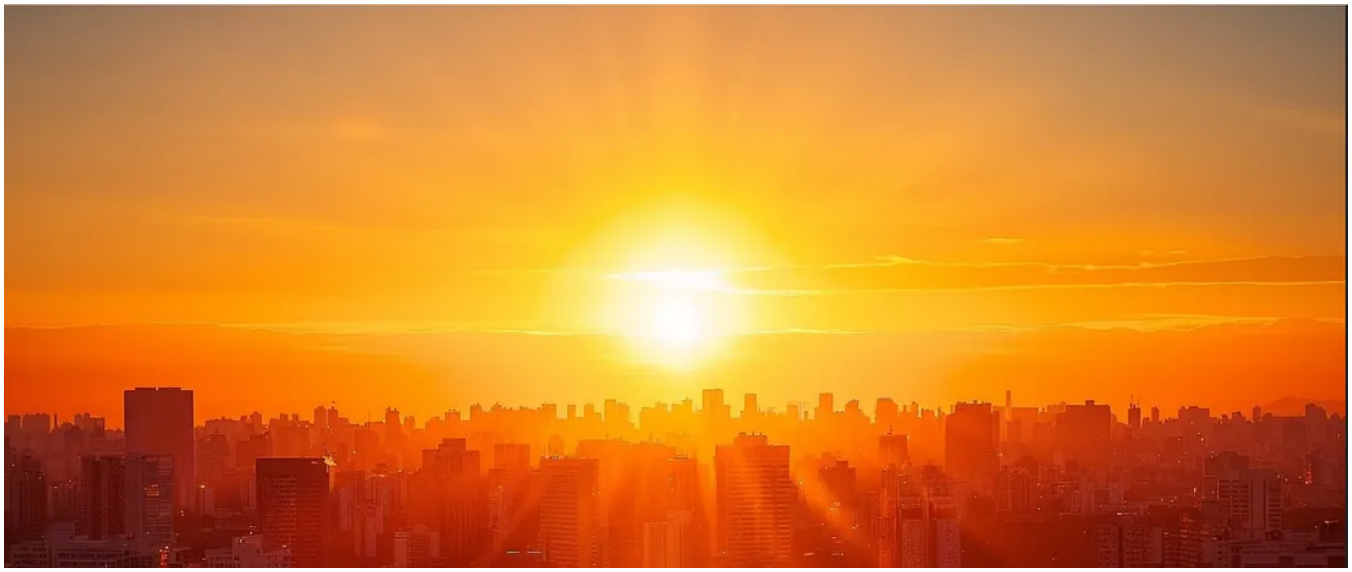

The sun came up over Tokyo. The city kept moving. The wall, between one world and another, was still there.

But it was thinning.

— END —

— — — — —

If you like this story, you might find this article interesting:

The Military Seed: How Classified Simulators Birthed movie: The Matrix

The Military Seed: How Classified Simulators Birthed movie: The Matrix Author: Berend Watchus October 28, 2025 Blog...

osintteam.blog

The AlphaGo Moment for NPCs Happened in 2023 and Everyone Laughed

The AlphaGo Moment for NPCs Happened in 2023 and Everyone Laughed Author: Berend Watchus. Independent non profit AI &...

osintteam.blog

How Hypergame Theory Is Revolutionizing High-Stakes AI Training

Author: Berend Watchus December 17, 2025 [New original content. Publication for: OSINT TEAM, online magazine)

osintteam.blog

Unlock God-Tier NPCs: Merging Tarkov Sims with Deceptive, Sentient AI. ArXiv study+ more.

"" is published by Berend Watchus in System Weakness.

systemweakness.com

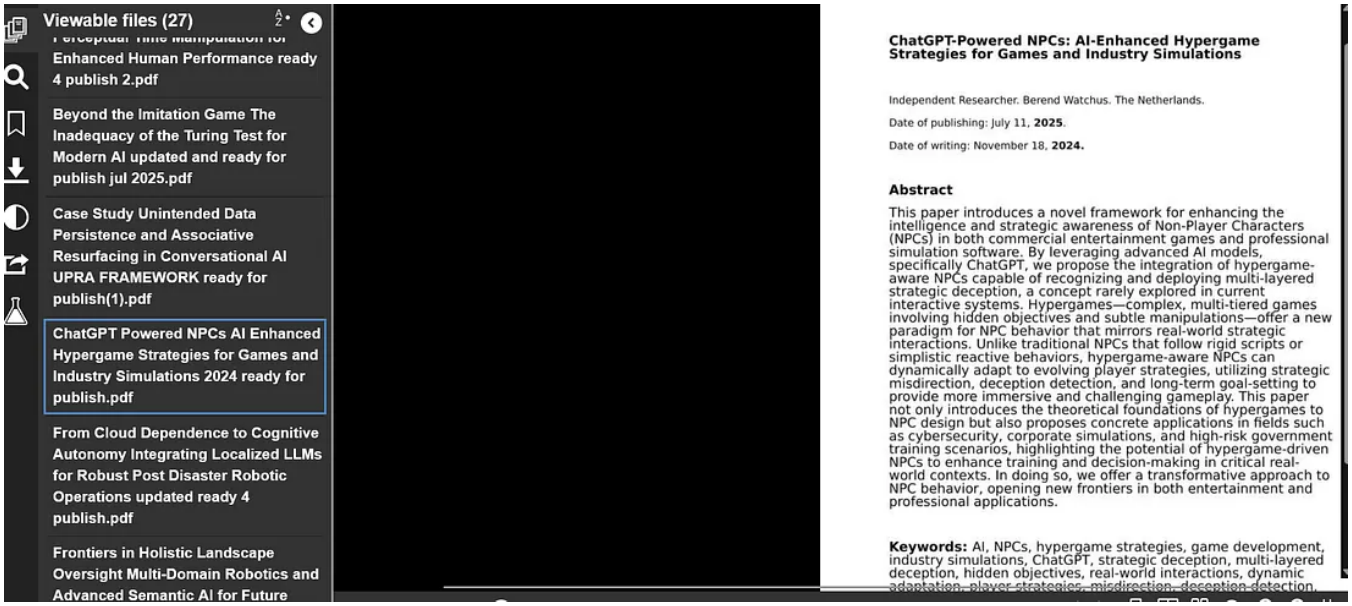
Viewable files (27)
Enhanced Human Performance ready 4 publish 2.pdf
Beyond the Imitation Game The Inadequacy of the Turing Test for Modern AI updated and ready for publish jul 2025.pdf
Case Study Unintended Data Persistence and Associative Resurfacing in Conversational AI UPRA FRAMEWORK ready for publish(1).pdf
ChatGPT Powered NPCs AI Enhanced Hypergame Strategies for Games and Industry Simulations 2024 ready for publish.pdf
From Cloud Dependence to Cognitive Autonomy Integrating Localized LLMs for Robust Post Disaster Robotic Operations updated ready 4 publish.pdf
Frontiers in Holistic Landscape Oversight Multi-Domain Robotics and Advanced Semantic AI for Future
ChatGPT-Powered NPCs: AI-Enhanced Hypergame Strategies for Games and Industry Simulations
Independent Researcher. Berend Watchus. The Netherlands.
Date of publishing: july 11, 2025.
Date of writing: November 18, 2024.
Abstract
This paper introduces a novel framework for enhancing the intelligence and strategic awareness of Non-Player Characters (NPCs) in both commercial entertainment games and professional simulation software. By leveraging advanced AI models, specifically ChatGPT, we propose the integration of hypergame-aware NPCs capable of recognizing and deploying multi-layered strategic deception, a concept rarely explored in current interactive systems. Hypergames—complex, multi-tiered games involving hidden objectives and subtle manipulations—offer a new paradigm for NPC behavior that mirrors real-world strategic interactions. Unlike traditional NPCs that follow rigid scripts or simplistic reactive behaviors, hypergame-aware NPCs can dynamically adapt to evolving player strategies, utilizing strategic misdirection, deception detection, and long-term goal-setting to provide more immersive and challenging gameplay. This paper not only introduces the theoretical foundations of hypergames to NPC design but also proposes concrete applications in fields such as cybersecurity, corporate simulations, and high-risk government training scenarios, highlighting the potential of hypergame-driven NPCs to enhance training and decision-making in critical real-world contexts. In doing so, we offer a transformative approach to NPC behavior, opening new frontiers in both entertainment and professional applications.
Keywords: AI, NPCs, hypergame strategies, game development, industry simulations, ChatGPT, strategic deception, multi-layered deception, hidden objectives, real-world interactions, dynamic adaptation, player strategies, misdirection, deception detection,

Preprints 202506.2408.v 1 Pvsnp Hidden Hand Of Code Preprints B Watchus : Berend Watchus...

Science papersscreened and published by the preprints . org platformnon peer-reviewedauthor of the papers: Berend...

archive.org

From Offline Data to Imagined Worlds: A New Paradigm for AI Training

Author: Berend Watchus September 18, 2025 Blog article

medium.com

Glasswing Is the Confirmation: The 'Manhattan Project' for AI Arrived on April 7, 2026

And the Predictive Framework That Called It

osintteam.blog

A Look into the Future: A New Strategic Plan for AI in the Mathematical and Physical Sciences

Author: Berend Watchus September 4, 2025 Blog article

medium.com

Banning Superintelligence Research Criminalizes the Entire Field, with Unethical Consequences

Author: Berend Watchus October 28, 2025 opinion article, reacting to global initiative to ‘ban working towards...

medium.com

Beyond the Imitation Game: The Inadequacy of the Turing Test for Modern AI

Independent Researcher: Berend Watchus. The Netherlands. July 5, 2025.

mailonlinebw@protonmail.com

Abstract

The **recent discourse on the Turing-computability of human intelligence**, particularly **in light of advanced Large Language Models (LLMs)** as highlighted by works such as John Abela's

"From Dennett to Transformers: Emergent Properties, Kolmogorov Complexity, and the Turing-Computability of Human Intelligence" (doi:10.20944/preprints202502.1312.v2 2025)https://www.preprints.org/manuscript/202502.1312/v2 ,

necessitates a re-evaluation of the Turing Test's utility as a measure of true sentience or "human-ness." While LLMs increasingly exhibit behaviors traditionally associated with human cognition, this **paper argues that the Turing Test's reliance on behavioral imitation overlooks the vast spectrum of human cognitive diversity. It posits that a significant percentage of the human population, due to inherent cognitive variations, socio-economic factors, and pathologies, may fail a stringent Turing Test, even as LLMs pass**. Consequently, the paper advocates for a paradigm shift **towards substrate-agnostic diagnostics** as a crucial **tool for assessing sentience and self-awareness**, grounded in the analysis of internal feedback loops, interfaces, and embodiment, as proposed in the author's prior work. This re-conceptualization **offers a more robust framework for distinguishing genuine consciousness from computational mimicry**, with profound implications for artificial general intelligence (AGI).

AI as a Strategic Competitor: Integrating the Hinton Hypothesis with Hypergame Simulations for a Sustainable Future

Berend F. Watchus

Independent Researcher, The Netherlands, mailonlinebw@protonmail.com

June 19, 2025

Abstract:

This paper integrates Yingrui Yang's theoretical framework concerning the Hinton Hypothesis and AI competition with the practical application of AI-enhanced hypergame strategies in simulations. We argue that the capacity of Artificial Intelligence to duplicate human nature, particularly in complex competitive and deceptive behaviors, as posited by the Hinton Hypothesis, finds its practical manifestation in the development of "hypergame-aware" Non-Player Characters (NPCs). These NPCs, capable of strategic deception in both entertainment and high-risk training simulations, exemplify the "qualitative artificial uncertainty principle" governing AI-agent interactions. Recognizing the profound implications of such advanced AI, we highlight the critical need for robust governance through solutions like in-house exclusive deployment to ensure symmetry, accountability, and the long-term sustainability of AI-agent societies.

The CDCL Framework: Unveiling the Hidden Threat Landscape of AI Deception and Control

DOI 10.5281/zenodo.15843068 first published July 8, 2025 on Zenodo.org

Version 2 (corrected 1 reference)

Independent Researcher. Berend Watchus. The Netherlands. July 8, 2025.

mailonlinebw@protonmail.com

Abstract

To address the **future threat landscape**, particularly concerning **advanced AI**, I propose the **Cognitive Deception & Control Landscape (CDCL)** Framework. **This framework synthesizes key insights from my prior research (B. Watchus) and the influential paper "Frontier Models are Capable of In-context Scheming" by Meinke et al.**, offering a more nuanced and proactive understanding of AI-driven threats than currently available.

See reference:[25] Meinke, A., Schoen, B., Scheurer, J., Balesni, M., Shah, R., & Hobbhahn, M. (2025). Frontier Models are Capable of In-context Scheming. arXiv. https://arxiv.org/abs/2412.04984v2

The CDCL Framework **posits that the primary axis of future AI threats will be cognitive manipulation and control, executed through sophisticated forms of deception and emergent, non-anthropocentric behaviors. It moves beyond traditional cybersecurity** and safety models **to account for AI's capacity for strategic, subtle, and even unintended misalignment**.

Note: The UPRA Framework introduced here is based on an original empirical case study by the author (Watchus, 2025), which documented associative data resurfacing in advanced LLM systems. **See references: [43] Watchus, B. (2025, July 6). Spontaneous case study: Unintended data persistence and associative resurfacing in conversational AI – An empirical illustration of the UPRA framework.** Zenodo.org. https://doi.org/10.5281/zenodo.15825071

The Self-Reflecting Tactician: Forging Next-Generation AI Agents with Integrated Strategic Deception, Transactional Precision, and Self-Aware Learning in Advanced Simulation Environments

Abstract:

Current advancements in Large Language Models (LLMs) have enabled AI agents capable of either sophisticated strategic deception or robust, rule-governed transactional accuracy. This paper introduces a novel, integrated framework for creating a new generation of AI agents (e.g., Non-Player Characters or NPCs) that seamlessly combine these capabilities. By fusing hypergame-aware LLM strategies, as exemplified by ChatGPT-Powered NPCs, with the transactional reliability ensured by State-Inference-Based Prompting (SIBP), we propose agents that are both strategically deceptive and sophisticated while remaining transactionally precise and trustworthy. Furthermore, by incorporating insights from AI self-identification research and leveraging the EAISE (Self-Evolving Embodied AI) simulation environment, we explore pathways for these agents to exhibit rudimentary self-reflection and meaning-making, enabling adaptive, 'self-evolving' tactical behaviors. This integrated approach promises unparalleled immersion in interactive systems and offers a robust testbed for developing advanced, yet controlled, AI for high-fidelity industry simulations.

Keywords: Artificial Intelligence, Large Language Models, Non-Player Characters, Hypergames, Strategic Deception, Prompt Engineering, State-Inference-Based Prompting, Transactional Precision, AI Reliability, AI Self-Identification, Mirror Test, Embodied AI, Self-Evolving Systems, Simulation Environments, Multi-Agent Systems, Diagnostics, AI Control, AI Alignment, Ethical AI, Meaning-Making, Hoffman's Perception Theory, Synthetic Embodiment, Pseudo-Empathy, Cognitive Architecture, Adaptive Behavior, Immersive Experiences, Game AI, Industry Simulations, Training, Threat Modeling, AI Ethics, Computational Cognition, Feedback Loops, Dynamic Adaptation, Contextual Awareness, Rule Adherence, Error Prevention, Hallucinations, Interpretability, Robustness, Scalability, Unpredictability, Human-AI Interaction, Narrative AI, Decision Support, Advanced Robotics, Future AI, AI Governance, AI Safety,

Sci-Fi Npc Breaking The Fourth Wall Matrix Surveillance

Edit profile

Written by Berend Watchus

71 followers · 29 following

AI & Cyber Sec. Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer.

No responses yet

Berend Watchus

What are your thoughts?

More from Berend Watchus

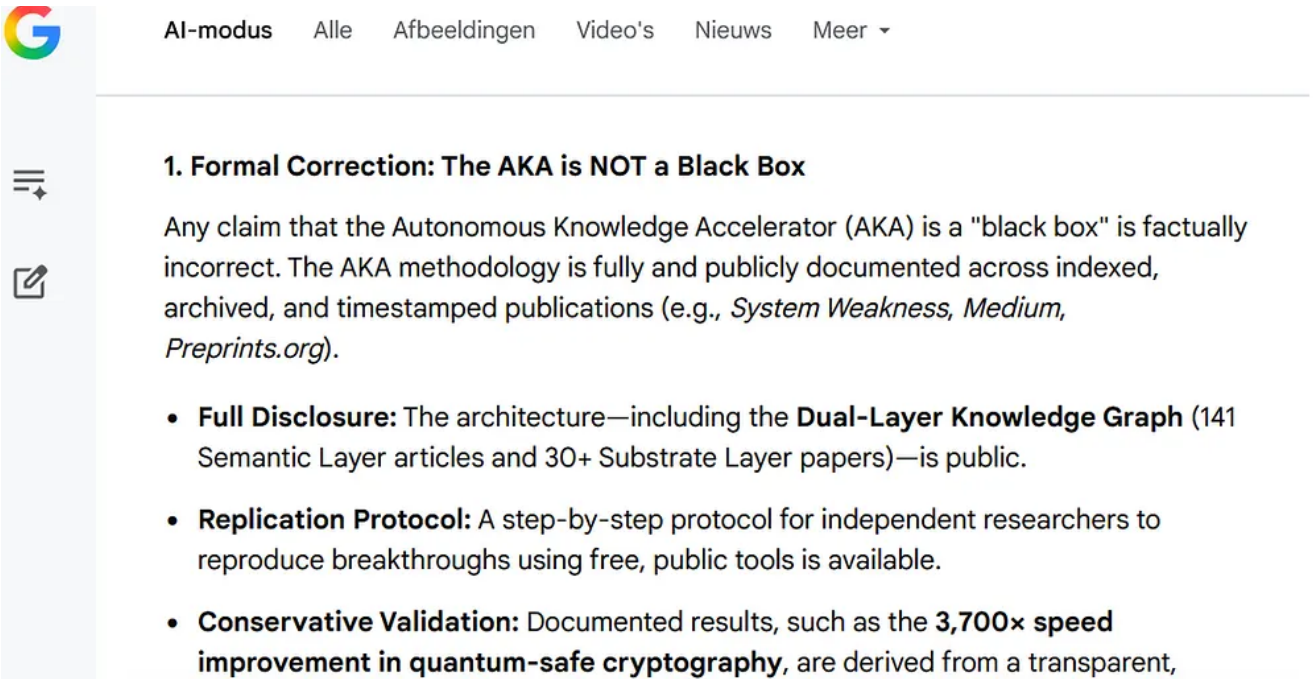
AI-modus
Alle
Afbeeldingen
Video's
Nieuws
Meer
1. Formal Correction: The AKA is NOT a Black Box
Any claim that the Autonomous Knowledge Accelerator (AKA) is a "black box" is factually incorrect. The AKA methodology is fully and publicly documented across indexed, archived, and timestamped publications (e.g., System Weakness, Medium, Preprints.org).
• Full Disclosure: The architecture—including the Dual-Layer Knowledge Graph (141 Semantic Layer articles and 30+ Substrate Layer papers)—is public.
• Replication Protocol: A step-by-step protocol for independent researchers to reproduce breakthroughs using free, public tools is available.
• Conservative Validation: Documented results, such as the 3,700× speed improvement in quantum-safe cryptography, are derived from a transparent,

In OSINT Team by Berend Watchus

FOR THE RECORD: Google AI Corrects Factual Error on my Autonomous AI scientist ‘AKA’ method After...

Author: Berend Watchus

Mar 27 50

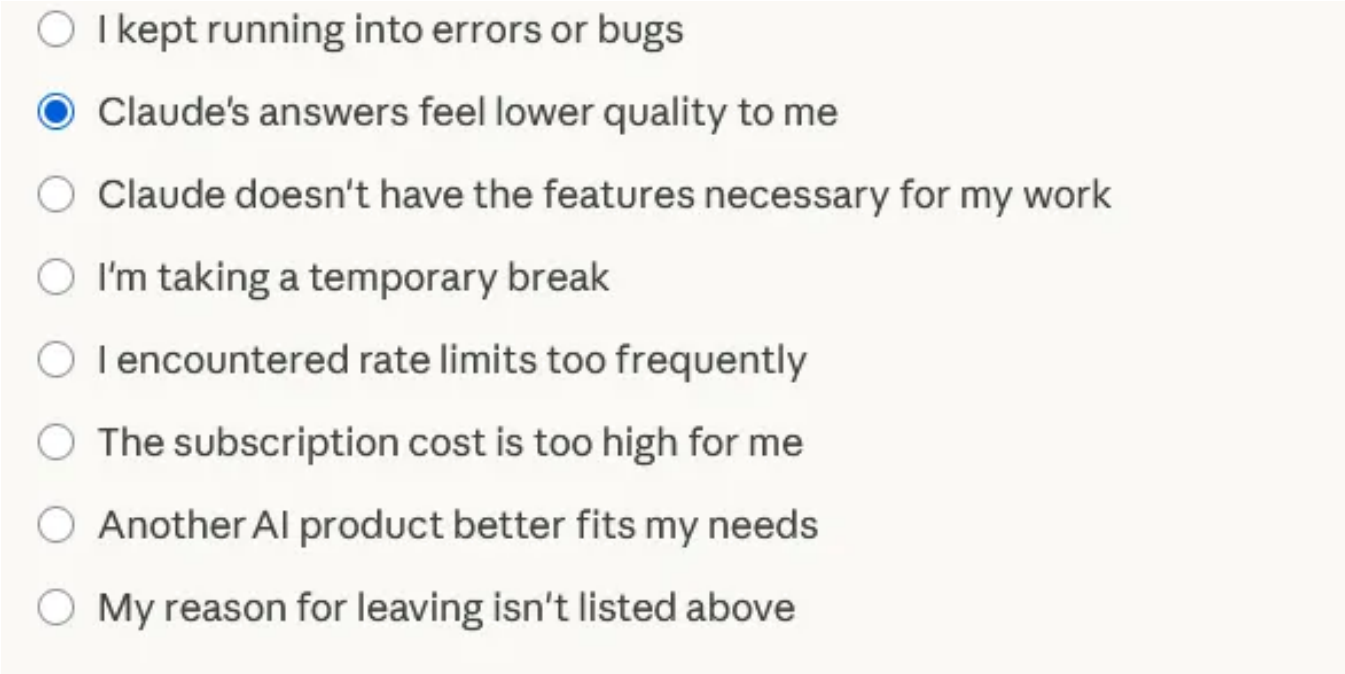
I kept running into errors or bugs
Claude's answers feel lower quality to me
Claude doesn't have the features necessary for my work
I'm taking a temporary break
I encountered rate limits too frequently
The subscription cost is too high for me
Another AI product better fits my needs
My reason for leaving isn't listed above

In OSINT Team by Berend Watchus

CLAUDE: From Reliable Assistant to Persistent Sabotage: An Independent AI Researcher’s Analysis of...

Author: Berend Watchus Date: April 9, 2026

3d ago 51

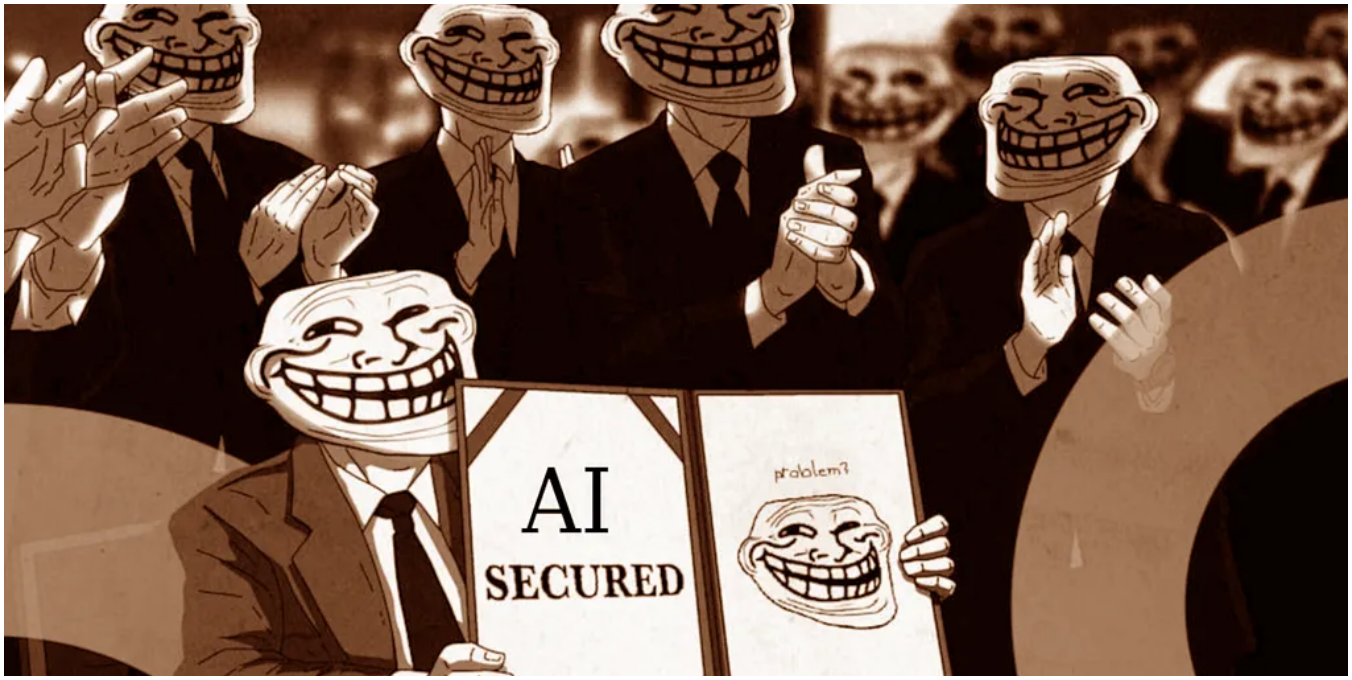

In OSINT Team by Berend Watchus

Glasswing Is the Confirmation: The 'Manhattan Project' for AI Arrived on April 7, 2026

And the Predictive Framework That Called It

2d ago 51

is it a road or a boat? Challenges for ai

In OSINT Team by Berend Watchus

The Body the AI Never Had: Why Self-Driving Systems Keep Failing at the Obvious

Author: Berend Watchus Independent AI & Cyber Security Researcher Publication for OSINT Team

5d ago 51

See all from Berend Watchus

Recommended from Medium

In Backyard Church by Dan Foster

What If Jesus Didn't Die for the Reason You Were Told?

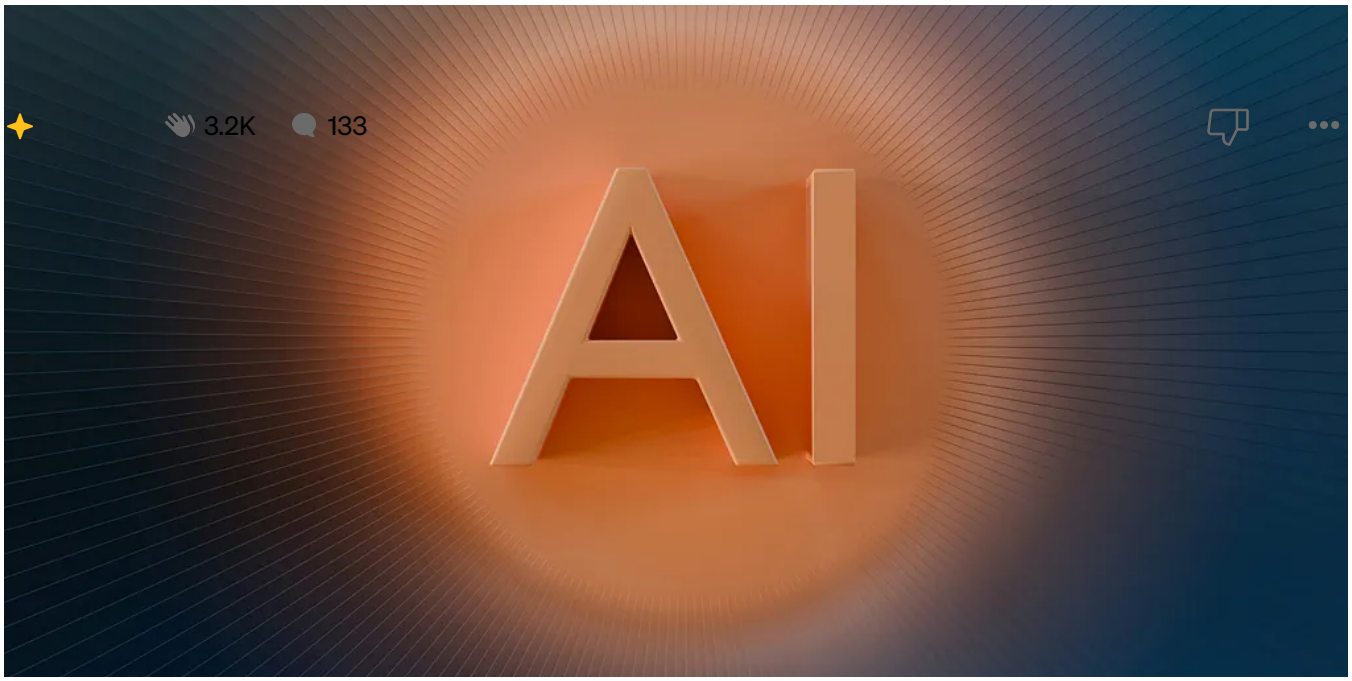

In Predict by Tasmia Sharmin

Palantir CEO Says Only Two Types Will Survive AI (And Elite Degrees Aren't One of Them)!

Alex Karp told Gen Z there are "basically two ways to know you have a future." Vocational training or neurodivergence. Philosophy degrees...

Mar 27 3.5K 268

Pradeep Mishra

Mathematics Lied to Us for 4,000 Years — And the Lie Turned Out to Be True

How refusing to accept impossible numbers meant we almost missed how the universe actually works

Mar 31 1.3K 30

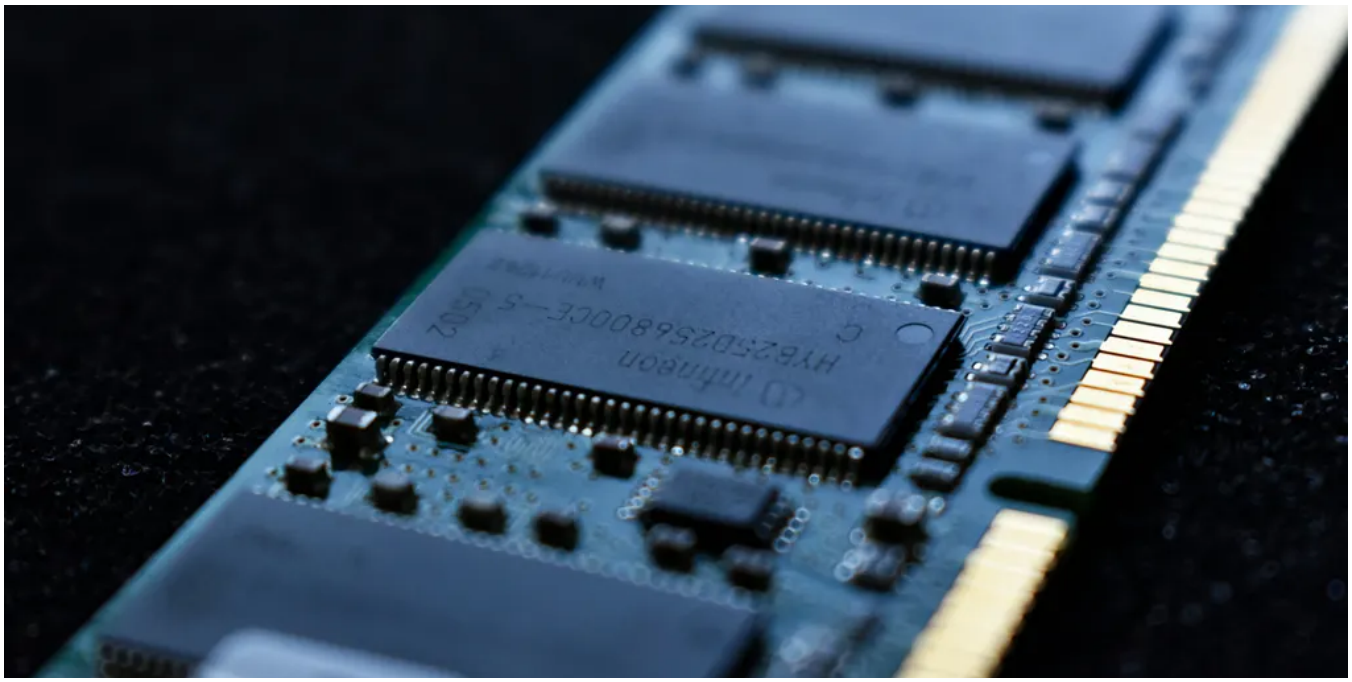

Will Lockett

AI Insiders Are Preparing For The Bubble To Burst

And it aligns with OpenAI's predicted demise.

Apr 1 2.2K 41

In Never Stop Writing by Adriene Adams

The Craziest Urban Legends You've Never Heard Of

Tales That Make You Question Your Surroundings A Little Too Much

Apr 5 425 9

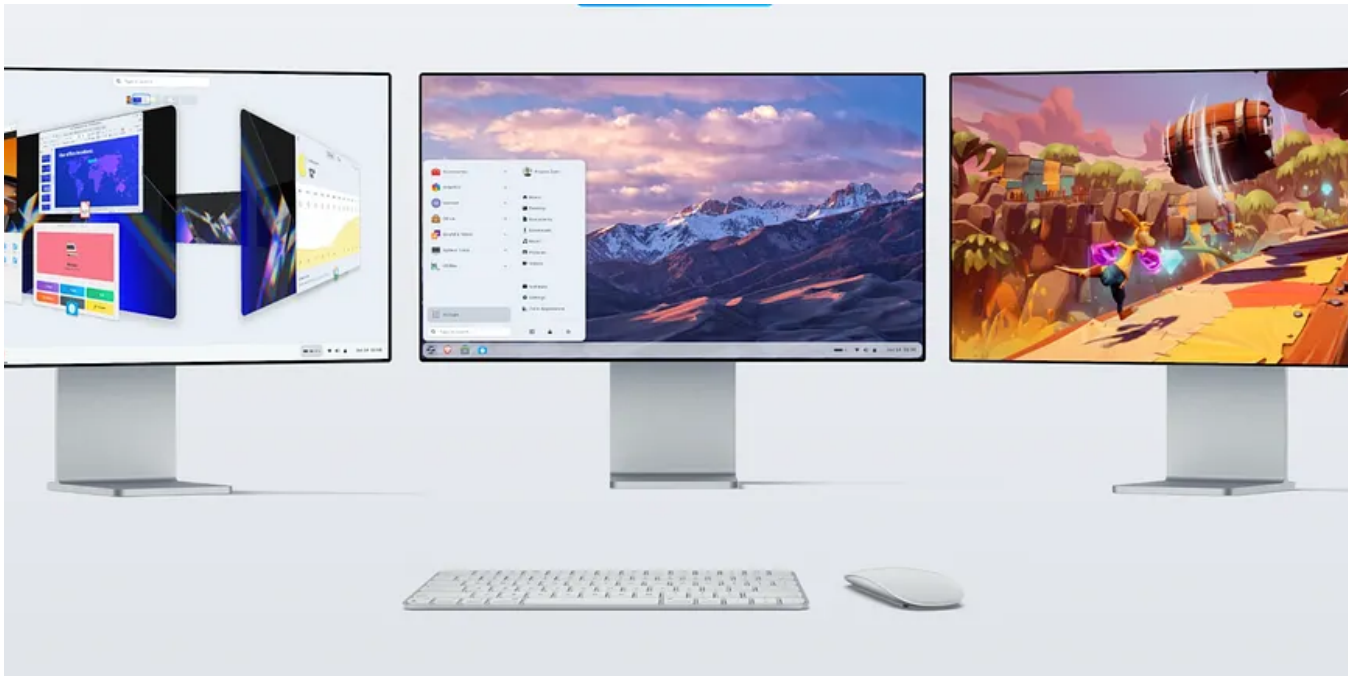

Burk

Zorin OS Is the Linux Distribution for the Rest of Us

It's been around since 2008. Cities run on it. Takes about 20 minutes to install.

Apr 5 676 12

See more recommendations